

unclassified

SECURITY CLASSIFICATION OF THIS PAGE

DTIC FILE COPY

②

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION <u>UNCLASSIFIED</u>			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution unlimited.		
JULE			5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR-TR-90-0075		
ER(S)			7a. NAME OF MONITORING ORGANIZATION AFOSR/NM		
6b. OFFICE SYMBOL (If applicable)			7b. ADDRESS (City, State and ZIP Code) AFOSR/NM, Bldg. 410 Bolling AFB DC 20332-6449		
6c. ADDRESS (City, State and ZIP Code) Raleigh, NC 27695-8205			8a. NAME OF FUNDING / SPONSORING ORG. AFOSR		
8b. OFFICE SYMBOL (If applicable) NM			9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER AFOSR-89-0124		
8c. ADDRESS (City, State and ZIP Code) AFOSR/NM, Bldg. 410 Bolling AFB DC 20332-6449			10. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO. <u>611647</u>	PROJECT NO. 3842	TASK NO. A5
11. Title (Include Security Classification) (DURIP) Two-Processor Alliant FX/4 System			12. PERSONAL AUTHOR(S) Dr. C. T. Kelley		
13a. TYPE OF REPORT Final		13b. TIME COVERED FROM 12/1/88 TO 11/30/89		14. DATE OF REPORT (Year, Month, Day) 1/16/90	
16. SUPPLEMENTARY NOTATION					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary; identify by block number)		
FIELD	GROUP	SUB-GROUP			
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
<p>This This proposal requested a two-processor Alliant FX/4 system with disk drives, tape backup, operator's console, and local line printer. This computer offers vector processing capability, a high-speed cache memory, and a shared memory multiprocessing architecture. When the project was funded, Alliant was willing to sell a two-processor FX/40 for the same amount.</p> <p>This system supports research in control and optimization, nonlinear least squares problems, partial differential equations, integral equations, parameter identification problems, and signal processing that is currently being carried out at North Carolina State University under AFOSR sponsorship by C. T. Kelley, R. J. Plemmons, M. Shearer, and S. J. Wright.</p> <p>The equipment is being used for design and testing of algorithms that take advantage of multiprocessing and vector architectures for large scale problems in these areas, for training of graduate students, and as a high-speed computing resource for problems, such as partial differential equations and large scale optimization problems that are too large to fit on the other local facilities presently available. Key words: compilation of research signal processing, control systems analysis, integral/least squares problems, control and estimation. (CMR)</p>					
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT			21. ABSTRACT SECURITY CLASSIFICATION		
<input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			<u>Unclassified</u>		
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Jon A. Sjogren			22b. TELEPHONE (Include Area Code) (202) 767-4940		22c. OFFICE SYMBOL NM

Final Technical Report AFOSR Grant #AFOSR-89-0124
Defense University Research Equipment

This is a summary of projects that have used or are using the Alliant FX/40 minisuper-computer purchased with funds from AFOSR grant #AFOSR-89-0124. The research title of the project was "(DURIP) Two-Processor Alliant FX/4 System". When the project was funded, Alliant was willing to sell as two-processor FX/40 for the same amount.

The proposal included four investigators, C. T. Kelley, R. J. Plemmons, M. Shearer, and S. J. Wright. The following summaries include titles, personnel, and a short description of the project. Some preprints of the work were attached to the semi-annual report. Preprints of work completed since the semi-annual report are attached.

Research Projects

Mesh Independence of Newton-like Methods for Infinite Dimensional Problems

Submitted for publication.

Preprint sent with semi-annual report.

C. T. Kelley and E. W. Sachs

This project is joint work between Professor E. W. Sachs, of the University of Trier in West Germany, and Kelley. Globally convergent modifications of Newton's method, such as the Armijo rule, can be applied to infinite dimensional problems and their discretizations. The results of this paper are that if the construction of the discretizations is done properly, then the convergence behavior of the iteration is the same for the discrete problems as it is for the infinite dimensional problem. Basic to these results is a new notion of convergence that is motivated by consideration of integral equations with continuous kernels. This result extends to the globally convergent case results of Allgower, Böhmer, Potra, and Rheinboldt, and the authors. In addition previous results on mesh independence of quasi-Newton methods were improved.

Numerical results were reported that illustrate the results. The Alliant was crucial to the success of the numerical experiments in this work. These experiments considered the performance of iterative methods for discretization of problems in function spaces, in this particular case, integral equations, as the discretization became finer. The largest problems reported on in the paper were dense, unstructured problems with 640 unknowns. In the course of the research itself, problems with as many as 6400 unknowns were solved. The algorithms considered in the paper have a natural parallel structure that could effectively exploit the Alliant's architecture. This structure is the topic of some work in progress by Kelley and a former Ph. D. student, J. I. Northrup, now at Worcester Polytechnic Institute.



By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

**A Fast Two-Grid Method for Matrix H-equations II:
Intergrid Transfers and Implementation**

Submitted for publication

Preprint enclosed

C. T. Kelley and J. I. Northrup

This project is joint work between Professor J. I. Northrup, of Worcester Polytechnic Institute and Kelley. This finished work represents a combination of joint work with Northrup on parallel methods for integral equations and Kelley's study of Nyström interpolation methods, two projects listed as in progress in the previous report. Northrup was a Ph. D. student of Kelley's and was supported by the AFOSR while a student.

In previous work of the authors quasi-Newton and multi-level algorithms for fully nonlinear integral equations were designed and analyzed. The motivating examples for that work were analogs of the Chandrasekhar H-equation for matrix-valued functions. In this paper we show how the performance of the algorithms can be improved by applying a quasi-Newton method to intergrid transfers and discuss implementation details on the Alliant FX series of multiprocessor computers.

**Fast algorithms for compact fixed point problems
with inexact function evaluations**

Submitted for publication.

Preprint enclosed.

C. T. Kelley and E. W. Sachs

This project is joint work between Professor E. W. Sachs, of the University of Trier in West Germany and Kelley. The project was listed as in progress in the semi-annual report as a joint project on parabolic boundary control problems with D. M. Hwang, a Ph. D. student of Kelley's who is supported by the AFOSR. Hwang's part has been split off and is described in a later abstract.

We describe and analyze a class of fast algorithms for computation of fixed points of completely continuous maps on Banach spaces. These algorithms are motivated by parabolic boundary control problems where the time integration is done by a high order backward difference formula, a variable stepsize, variable order method, or a combination of such methods. In these cases, the nonlinear maps do not have the smoothness or collective compactness properties required by known fast algorithms. We show how a multi-level technique of Atkinson can be modified to attack such problems, discuss how quasi-Newton methods can improve performance, and show how our approach can be applied to parabolic boundary control problems.

The Alliant is being used to implement our algorithms, using L. Petzold's code, DASSL, as the time integrator. Our other computing environment, a SUN workstation, is far too slow to permit the testing and comparison of algorithms that we must do.

Applications of Broyden's Method in Banach Spaces

In Progress

C. T. Kelley and D. M. Hwang

Superlinear convergence results are given for the first time for Broyden's method in non-Hilbert Banach spaces. Numerical observations on the Alliant illustrate the dependence of the superlinear rates on the topology.

The Harmonic Balance Method in Computer Aided Design of Microwave Devices

In Progress

C. T. Kelley, R. J. Trew, P. Gilmore, L. Mukundan, and D. E. Stoneking

This project is joint with Professor R. J. Trew of the Electrical and Computer Engineering department at North Carolina State University, D. E. Stoneking, a Ph. D. student of Trew's, L. Mukundan and P. Gilmore, Ph. D. students of Kelley's, and Kelley. The harmonic balance method is a projection method for a nonlinear two point boundary value problem for an integro-differential equation that arises in certain device models used in computer aided design. Kelley and his students are analyzing convergence properties of this projection method and designing optimization algorithms that take advantage of the structure of the particular harmonic balance problems that Trew and his group encounter in their modeling code, TEF-LON. The important characteristics of these nonlinear problems are that function evaluations are extremely expensive and sequences of problems often need to be solved. This means that quasi-Newton and continuation methods should be used to reduce cost. The standard quasi-Newton method for dense unstructured problems, Broyden's method, while better than no derivative updating at all, does not provide superlinear convergence in the limit of infinitely fine discretization. The ultimate goal of this research is to design a new type of quasi-Newton method that has better convergence properties and implement it in a globally convergent way. We currently use Broyden's method and recompute the Jacobian when the Broyden direction is not a descent direction and are beginning to implement an Euler-Newton continuation scheme. Intensive testing of these algorithms will be possible when TEF-LON has been ported to the Alliant. This port is now in progress, and we are discovering opportunities for exploitation of the parallel/vector architecture of the Alliant.

Hydrodynamic Hot Electron Transport Simulation based on the Monte Carlo Method

Submitted for publication

Preprint enclosed with semi-annual report.

C. T. Kelley, D. L. Woolard, J-L. Pelourad, R. J. Trew, and M. A. Littlejohn

A hydrodynamic hot electron model is used to study electron transport through a sub-

micron $N^+ - N - N^+$ GaAs structure. This study is used to investigate improvements which the unique features of this model offer to analysis of devices operating under nonstationary transport conditions. The model is based upon semiclassical "hydrodynamic" conservation equations for the average carrier density, momentum and energy. The general model includes particle relaxation times, momentum relaxation times, energy relaxation times, electron temperature tensors and heat flow vectors as a function of average carrier energy for the Γ , X and L valleys of GaAs. For this study, we utilized a simplified single electron gas version of our model to clearly reveal the impact of the nonstationary terms in the model. Results from both a drift-diffusion model approach and a Monte Carlo analysis are used to show the relative accuracy and facility this new model offers for investigating practical submicron device structures operating under realistic conditions.

**Simulation of the Variation and Sensitivity of GaAs MESFET
Large-Signal Figures-of-Merit due to
Process, Material, Parasitic, and Bias Parameters**

Submitted for Publication

Preprint Enclosed

D.E. Stoneking, R.J. Trew, and L. Mukundan

A simulator for calculating MESFET large-signal figures-of-merit and their sensitivities with respect to various device design, material, and operational parameters has been developed. A study based upon an ion-implanted device with a $0.42\ \mu\text{m}$ gate length and $1.0\ \text{mm}$ gate width is presented. The figures-of-merit of the study are maximum power-added efficiency, output power at the maximum power-added efficiency, and output power at 1 dB gain compression. The study parameters are peak implant doping, implant straggle, implant range, gate length, gate width, and device gate-drain breakdown voltage.

Mukundan is a Ph. D. student of Kelley's. This work represents use of the Alliant for scientific computing in an engineering setting.

Recursive Least Squares Computations

Submitted for Publication

Preprint Enclosed

Robert J. Plemmons

We consider parallel implementations of algorithms for recursive least squares computations based upon the information matrix and the covariance matrix updating methods. The target architecture is a shared-memory multiprocessor, and test results on an Alliant system with 2 vector processors demonstrate the parallel efficiencies of the algorithms. The results also show that the covariance method in a form suggested by Pan and Plemmons is easily the most efficient on the Alliant multiprocessor. Applications include robust regression in statistics and modification of the Hessian matrix in optimization, but the primary motiva-

tion for this work is the need for fast algorithms for recursive least squares computations in signal processing.

Implicit Nullspace Iterative Methods for Constrained Least Squares Problems

Submitted for publication

Preprint enclosed

Douglas James

We propose a class of iterative algorithms for solving equality constrained least squares problems, generalizing an order-reducing algorithm first analyzed by Barlow, Nichols, and Plemmons. These algorithms, which we call implicit null space methods, are based on the classical nullspace method, except that a basis for the nullspace of the constraint matrix is not explicitly formed. The implicit basis acts as a preconditioner for a set of (unformed) normal equations. The methods allow great flexibility in the choice of preconditioner, and are suitable for parallel implementation on substructured problems. We offer some numerical results for both structural engineering applications and Stokes Flow.

Order-Reducing Conjugate Gradients versus Block AOR for Constrained Least Squares Problems

Submitted for publication

Preprint enclosed

Douglas James

We compare the convergence properties of two iterative algorithms for solving equality constrained least squares problems. The first algorithm, due to Barlow, Nichols, and Plemmons, applies a variation of the conjugate gradient algorithm to a symmetric positive definite system which is smaller than the original problem. The second, Block Accelerated Over-relaxation, is a two parameter generalization of block SOR. Barlow, Nichols, and Plemmons have proven that their order-reducing conjugate gradient algorithm converges faster than block SOR. We extend their result to show that the algorithm is also superior to block AOR. Numerical examples confirm the analysis.

An Iterative Substructuring Algorithm for Equilibrium Equations

In Progress

Preprint enclosed with semi-annual report.

R. J. Plemmons and Douglas James

This is joint work between Plemmons and a Ph. D. student, Douglas James. The topic of iterative substructuring methods, and more generally domain decomposition methods, has been extensively studied over the past few years, and the topic is well advanced with respect

to first and second order elliptic problems. However, relatively little work has been done with regard to application to general equilibrium equations such as those arising, for example, in realistic structural analysis problems. The potential for effective use of iterative algorithms here is good, but such methods are still far from being competitive with direct methods in industrial codes. The purpose of this paper is to investigate a preconditioned conjugate gradient method for the Kuhn-Tucker equations associated with constrained least squares, suggested by Barlow, Nichols and Plemmons, in the context of substructuring methods. We propose to use a mixed approach, consisting of both direct reduction in the substructures and the conjugate gradient based iterative algorithm to complete the computations. Some computational experience on an Alliant FX/40 vector multiprocessor with the algorithm applied to a structures problem and a fluid flow problem gives an indication of the efficiency of our approach.

The use of an Alliant FX/40 vector purchased under a grant from the AFOSR is proving essential to the timely completion of this work, which is still in progress. In particular the vector-multiprocessor capabilities of this computer is highly suited to the matrix \times vector products (BLAS 2) operations involved in the conjugate gradient based iterative algorithms we are implementing.

Incomplete QR Factorizations for Sparse Unstructured LS Problems

In Progress

Preprint enclosed with semi-annual report.

R. J. Plemmons and Douglas James

We are studying the iterative solution of large sparse unstructured LS problems involving a coefficient matrix A of full column rank. Our strategy involves performing an approximate QR factorization of A , and using the resulting R (which is the Cholesky factor of $A'A$) as a preconditioner for the conjugate gradient algorithm applied to the factored form of the normal equations. We form R using Givens rotations, retaining only some of the non-zero entries produced by the rotations. We have established that such incomplete QR factorizations can break down (producing a singular R) given virtually any strategy for retaining non-zeros, even under very restrictive conditions on A and $A'A$. We've conducted experiments on the FX/40 (using matrices from the Boeing-Harwell test collection) showing that such breakdowns do in fact occur in practice. Given any of several fast automatic reordering strategies, however, we can eliminate the breakdowns on the test problems, and produce a fairly effective preconditioner (in terms of iterations required for convergence). Unfortunately, however, the total execution time for the algorithm is not yet competitive. We are currently studying ways to improve the quality of the preconditioner produced by this process.

Implementing Proximal Point Methods for Linear Programming

To appear in Journal of Optimization Theory and Applications

Preprint enclosed with semi-annual report

Stephen Wright

Aspects of the efficient implementation of proximal point methods for large-scale linear programs have been investigated. The proximal point algorithm is a very general technique which can be applied to a large class of constrained optimization problems. The well-known method of multipliers is a special instance of it. The application of the latter method to linear programming problem has been investigated previously by Mangasarian and several Soviet authors. Our aim in this project is to investigate other instances of the proximal point algorithm in this context, and in particular, to find efficient algorithms for solving the subproblem which occurs at each main iteration, namely, a convex quadratic program with simple non-negativity constraints. We have obtained numerical results from Alliant and CRAY implementations of the resulting algorithms. A two-phase algorithm, which works by using the proximal point approach in the first phase, followed by the simplex code MINOS in the second phase, has given superior performance to MINOS alone on some randomly generated problems. Parallel implementation on the Alliant has been investigated. The central issue is one of parallelizing various sparse matrix-vector operations, and good efficiency has been obtained.

A Parallel Algorithm for Banded Linear Systems

Submitted for publication

Preprint enclosed with semi-annual report

Stephen Wright

A partitioned Gaussian elimination algorithm with partial pivoting which is suitable for multiprocessors with small to moderate numbers of processing elements has been developed. We only assume that the system is non-singular; hence the submatrices in our chosen partitioning may be rank-deficient, and this makes the algorithm more complex than those which have been proposed for diagonal-dominant and symmetric positive definite systems. We have examined the effect on the solution accuracy of ill-conditioning in the submatrices. Numerical results have been obtained on Alliant, Sequent and Encore multiprocessors, with the help of the SCHEDULE parallel programming package which has been developed at Argonne National Laboratory. Ongoing research is centering on applications of the resulting algorithm to discrete optimal control problems. Preliminary results indicate that substantial speedup over the traditional recurrence relations is possible in a multiprocessor environment.

Solution of Discrete-Time Optimal Control Problems on Parallel Computers

Submitted for publication

Preprint enclosed

Stephen Wright

We have investigated locally-convergent algorithms for discrete-time optimal control problems which are amenable to multiprocessor implementation. Parallelism is achieved both through concurrent evaluation of the component functions and their derivatives, and through the use of a parallel band solver which solves a linear system to find the step at each iteration. Results from an implementation on the Alliant are described.

Algorithms for Minimization Subject to Bounds

Submitted for publication

Preprint enclosed with semi-annual report

Stephen Wright

We have investigated algorithms for minimization of functions of many variables subject to bound constraints on those variables. The main algorithm uses a trust region strategy for minimizing the corresponding exact ℓ_1 -penalty function. The subproblem contains component-wise nonsmooth terms, and is solved using a method similar to two-metric gradient projection methods for constrained optimization. This algorithm is compared to others which (i) minimize the ℓ_1 -penalty function directly without using trust regions, (ii) use two-metric gradient projection with line search, and (iii) apply a trust region method and take the constraints into account explicitly. It is found that the trust-region methods tend to require fewer function and gradient evaluations, at the cost of increased matrix manipulation during each (outer) iteration. Numerical results have been obtained on the Alliant.